

# Automatic Target Recognition Directed Image Compression<sup>†</sup>

Charles D. Creusere

Alan Van Nevel

Naval Air Warfare Center, China Lake, CA 93555

chuck@wavelet.chinalake.navy.mil

alan@peewee.chinalake.navy.mil

## Abstract

*In this paper, we present a novel synthesis of two separate areas of image processing: automatic target recognition/cueing (ATR/ATC) and embedded image compression. In order to maximize the information content of a transmitted image, an ATR algorithm is used to detect potential areas of interest, and the compression algorithm regionally compresses the image, allocating more bits to the areas of interest. In this fashion, contextual information is retained, albeit at a lower resolution. Examples of the hybrid algorithm are presented.*

## 1. Introduction and Background

As the number of operational UAV's (unmanned aerial vehicles) and UCAV's (unmanned aerial combat vehicles) increases, the availability of sufficient communications bandwidth will become a major concern. While command and control signals require a small portion of the available bandwidth, high resolution imagery and video can easily consume all of the available bandwidth and more. Unfortunately, tactical datalinks such as Link 16 or SATCOM only support data rates of 57 kilobits/s (with error correction) and 4.8 kilobits/s respectively, while other high speed systems like CDL (common datalink) have significant limitations on the number of simultaneous users allowed. To increase the number of systems that can operate simultaneously using the limited RF bandwidth available, one must adopt digital compression. However compression results in degraded imagery at the receiver for reasonable compression ratios (>2:1). With these severe limitations on the bandwidth in mind, we have developed a hybrid algorithm which draws upon ATR/ATC algorithms and regional embedded compression techniques.

The basic idea of the algorithm is relatively simple. While an image may contain several different objects (trucks, buildings, tanks, etc.) only a few may be of

potential interest to a military observer. An ATR algorithm is used to select potential areas of interest, which are then fed to a compression algorithm, that regionally compresses the image, allocating more bits to the areas of interest. This process results in an image which has areas of high resolution (potential targets) and areas of low resolution in which contextual information is retained. Overall, a high compression ratio can be achieved while maximizing information content. In the first section, we will describe the ATR algorithm, developed by Mahalanobis *et al* [1-4], which we use to determine the regions of interest. In the following sections, we will describe the embedded zerotree wavelet (EZW) and how it has been adapted for spatially variant resolution. In the last sections we present some results and discuss potential applications and future directions.

## 2. MACH Filters

### 2.1 Mathematical Overview

In order to maximize the information content sent over the available bandwidth, we have chosen to identify potential regions of interest. To accomplish this, we have implemented a class of correlation filters as developed by Mahalanobis *et al*. These filters have exceptional tolerance to scaling and rotation distortions. The tolerance of the filters is incorporated through the selection of an appropriate training set, and can be tuned to provide high (generalization) or low (specificity) tolerance.

In the discussion of the MACH (maximum average correlation height) filters that follows, bold lowercase indicates a column vector, while bold uppercase represents a diagonal matrix. The filters result from maximizing the ratio

$$J(\mathbf{h}) = \frac{|\mathbf{h}^+ \mathbf{m}|}{\mathbf{h}^+ \mathbf{S} \mathbf{h}} \quad (1)$$

<sup>†</sup>approved for public release; distribution unlimited  
this paper is declared a work of the U.S. Government and is not subject to  
copyright protection in the United States.

19981123 136

where  $\mathbf{h}$  is the correlation filter and  $\mathbf{m}$  is the average of the training images in the Fourier domain. Each image is lexicographically ordered to form a vector.  $\mathbf{S}$  is the average similarity measure matrix

$$\mathbf{S} = \sum_{k=1}^N (\mathbf{X}_k - \mathbf{M})(\mathbf{X}_k - \mathbf{M})^+ \quad (2)$$

In eq. (2)  $\mathbf{X}_k$  are the individual training images, again in the Fourier domain. The training image is lexicographically ordered and its elements placed on the diagonal of  $\mathbf{X}_k$  while  $\mathbf{M}$  is the mean training image, arranged similarly to  $\mathbf{X}_k$ . Furthermore, all of the processing to generate the filters is performed in the Fourier domain to gain translational invariance. It is possible to perform the processing in other domains (e.g. wavelet or spatial) but care must be taken to properly register the training imagery.

The optimal filter  $\mathbf{h}$  is then given by

$$\mathbf{h} = \mathbf{S}^{-1} \mathbf{m} \quad (3)$$

Variants on the MACH filter can be achieved by varying the performance metric one wishes to maximize. For example Refrieger [5] has developed optimal trade-off synthetic discriminant filters (OTSDF's) which attempt to minimize the energy functional

$$E(\mathbf{h}) = \mathbf{h}^+ \mathbf{Q} \mathbf{h} - \delta |\mathbf{h}^+ \mathbf{m}| \quad (4)$$

where

$$\mathbf{Q} = \alpha \mathbf{P} + \beta \mathbf{D} + \gamma \mathbf{S} \quad (5)$$

$\mathbf{S}$  is as defined previously,  $\mathbf{P}$  is the power spectral density of the expected noise, and  $\mathbf{D}$  is the average power spectral density of the training set. The constants  $\alpha, \beta, \gamma, \delta$  are non-negative and must satisfy  $\alpha^2 + \beta^2 + \gamma^2 + \delta^2 = k$  where  $k$  is any positive constant. Minimizing  $E(\mathbf{h})$  results in

$$\mathbf{h} = \frac{\delta}{2} \mathbf{Q}^{-1} \mathbf{m} \quad (6)$$

By varying the parameters, one can optimize filter performance for the situation under study. If one sets  $\alpha = \beta = 0$ , the result is the MACH filter discussed earlier. Further variations can be made to the basic idea, including the extension to multiple class discrimination using distance classifier correlation filters (DCCF's), which are able to distinguish between multiple classes of similar objects (e.g. T72's vs. M1A1 tanks).

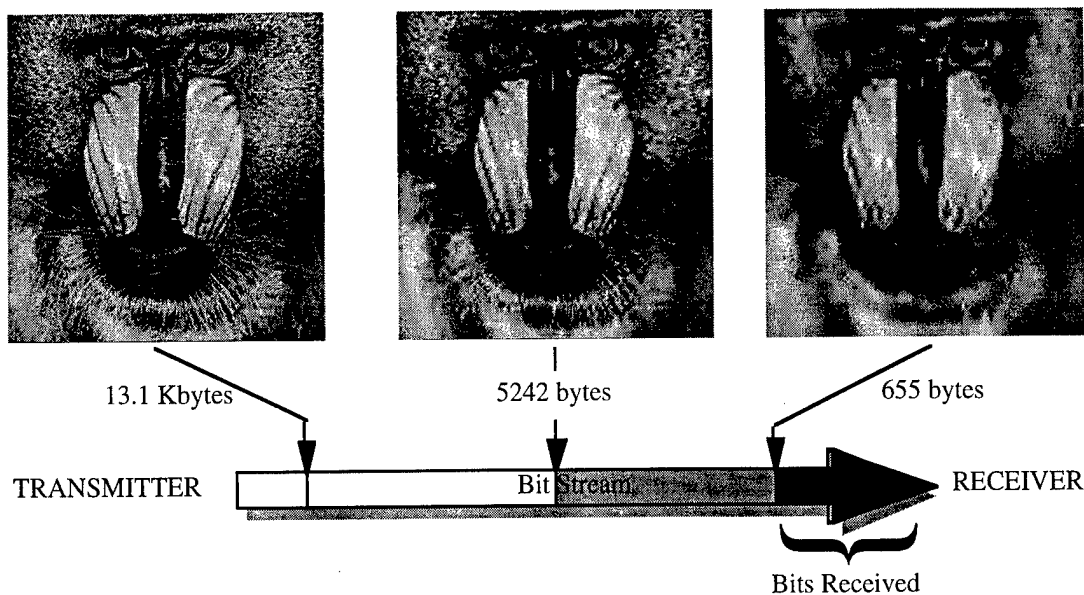
The class of MACH filters was chosen for the feature detection for several reasons. As discussed, the filters can incorporate varying degrees of distortion tolerance and can be built to generalize classes of targets. Another benefit of the algorithm is that the result is statistically optimum and depends on a realistic, mathematically rigorous optimization procedure as opposed to other heuristic methods. A final consideration is the computational efficiency. The MACH filters require no segmentation or edge detection preprocessing and the correlation step can be performed rapidly using dedicated FFT hardware.

## 2.2 MACH implementation

To implement the MACH filters, one must first decide upon a representative training set. Typically, the training set consists of  $N < 20$  images from varying perspectives. A training set of one image will result in a filter similar to the matched filter with no distortion tolerance while having dozens of perspectives and scalings will produce a filter with a broad response and low discrimination properties. The filter  $\mathbf{h}$  is first calculated off-line from the training data. If one is using the OTSDF's, some parameter tuning can be done at this point to maximize the correlation peaks for the training data.

Following correlation of an input test scene with  $\mathbf{h}$ , the correlation scene must be processed to determine the areas of interest. Previous correlation filters [6-8] had placed constraints on the correlation height, and classification was then accomplished by comparing the correlation height of the test scenes to the constraint. Generally, when using the correlation height as a metric for detection and/or classification, a threshold must be set. By changing this threshold one can trade off between the probability of detection and the probability of false alarms, a lower threshold allowing more false alarms and a higher threshold reducing the probability of detection.

A second metric that can be used involves comparing a local correlation energy to the global correlation energy. A square window (1x1, 3x3, 5x5...) is chosen and centered around the correlation peak. The energy within this window is calculated and the ratio of the local energy to the global energy is calculated. This ratio can then be compared to a threshold for detection and classification. The local energy percentage can be modified to allow for multiple target possibilities by selecting multiple windows based on correlation height and excluding these energies from the global energy calculation. The benefit of this approach is that the ratio is independent of illumination or amplification



**Figure 1.** Embedded image compression.


effects. The overall peak height can be affected by constant amplification but the ratio will remove this problem. This metric works well in rejecting false peaks due to clutter since most correlation surfaces for clutter images will not contain a high percentage of energy in a localized window.

In our hybrid algorithm, the second metric was chosen to determine regions of interest. No hard thresholding was used for detection. Instead, the top three energy percentage locations were selected as potential regions of interest to be compressed at a higher resolution than the background. The choice of three targets is somewhat arbitrary and can be changed based on the application. If a large number of areas are desired at high resolution, it may impact how the coding of the side information is performed. The choice may also be eliminated completely with the use of thresholding to eliminate false alarms and to increase the probability of detection. In this demonstration it was sufficient to designate a number of potential targets to effectively illustrate the concept.

### 3. Embedded Zerotree Wavelet Compression

At the core of our feature-based approach to compression is an embedded coding algorithm. In this method of compression, data is transmitted to the receiver in order of importance-- i.e., that data which most reduces the error between the reconstructed image and the original image is sent first. This concept is illustrated in Fig. 1. There are a number of advantages to embedded compression algorithms: fixed bit rates (e.g., compression ratios) are easily achieved, unequal


transmission error protection is trivial, and inherently robust bit streams can be created.



**Figure 2.** Wavelet coefficient mapping with one complete zerotree shown. Note that the wavelet scale ( $s_0$ ,  $s_1$ , etc.) is inversely proportional to the spatial frequency.

The fundamental observation which inspired the EZW algorithm [9] is that there is a strong correlation between insignificant coefficients at the same spatial locations in different wavelet scales-- i.e., if a wavelet coefficient at a coarser scale is zero, then it is more likely that the corresponding wavelet coefficients at finer scales will also be zero. Figure 2 shows a 3-level, 2D wavelet decomposition and the links which define a


single zerotree (the quadtree data structure containing all of the coefficients corresponding to a given region of the original image). If a wavelet coefficient at a given scale is zero along with all of its descendants (as shown in Fig. 2), then a special zerotree-root symbol (ZTR) is transmitted, eliminating the need to transmit the values of the descendants. Note that ZTR symbols can be created at any level of the wavelet coefficient mapping. Thus, if one of the three zerotree children at level  $s_2$  in Fig. 2 is significant for a given bit plane while the others are not, then the low frequency root coefficient must be represented by an isolated zero (IZ) symbol (assuming that it is also insignificant) while the two insignificant children are represented with ZTR symbols. Basically, one can view this redundancy extraction process as a multiresolutional, self-terminating run-length encoder. Regardless of how one views it, zerotrees exploit the correlation of insignificance across scales and decrease the number of symbols which must be processed by the arithmetic encoder and transmitted to the receiver.



**Figure 3:** Scanning order of wavelet coefficients.

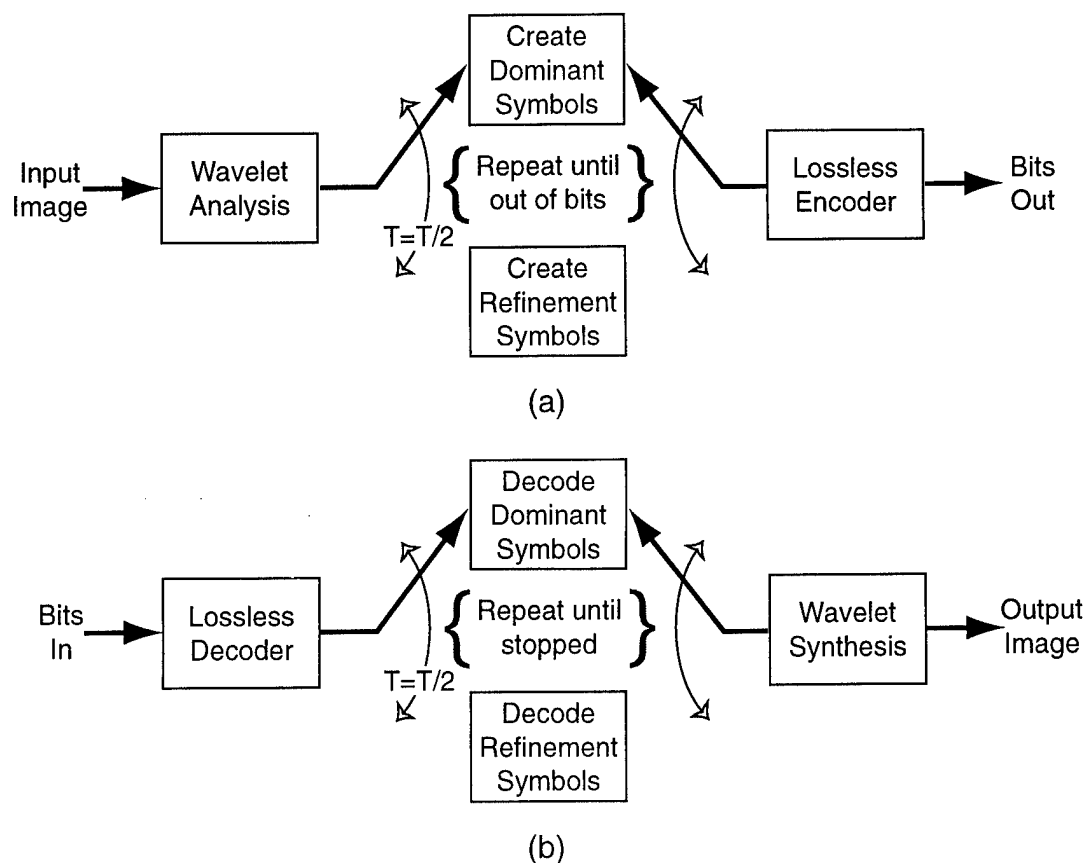
In order to generate an embedded code (where information is transmitted in order of importance), the algorithm scans the wavelet coefficients in what is basically a bit-plane fashion. First, a starting threshold is selected that is at least  $1/2$  as large as the magnitude of the largest wavelet coefficient. If the starting threshold is selected to be a power of 2, then a very fast approach can be used to compute all of the zerotree dependencies in one pass through the wavelet coefficients [10]. Starting with the appropriate threshold, the algorithm sweeps through the coefficients from low to high frequency subband as shown in Fig. 3, transmit-

ting the sign (+ or -) if a coefficient's magnitude is greater than the threshold (i.e., it is significant), a ZTR if it is less than the threshold and the root of a zerotree at the coarsest possible scale, or an IZ otherwise— this is the dominant pass. Next for the subordinate pass all coefficients deemed significant in the dominant pass are added to a second subordinate list which is itself scanned. One bit is transmitted for each coefficient on this list during the pass, decreasing its approximation error in the decoder by  $1/2$  (the coefficient's absolute error during a given pass depends on the value of the starting threshold). One iteration of this successive refinement process is illustrated by Fig. 4. The threshold is then halved and the two passes are repeated with those coefficients having been previously found significant being replaced by zeros in the dominant pass (so that they do not inhibit the formation of future zerotrees). The symbol stream created by this scanning process is then passed through an arithmetic encoder to eliminate any remaining statistical redundancy before transmission to the decoder. A block diagram of the complete process is shown in Fig. 5a. To estimate symbol probabilities for the arithmetic encoder and decoder, we use the simple single-context, backward-adaptive model presented by Witten, *et al.* in [11]. Slightly better compression can be achieved using the multicontext model proposed by Shapiro at the cost of decreased execution speed. The routine of dominant pass, subordinate pass, and threshold reduction continues until the bit budget is exhausted or until some distortion criterion is reached; at that point, the encoder transmits a stop symbol and begins processing the next frame in the video sequence.



**Figure 4.** One iteration of the successive refinement process.

The image decoder simply inverts each operation performed by the encoder in reverse order: i.e., it arithmetically decodes the bit stream to create symbols and then it decodes the symbols to progressively refine its estimates of the wavelet coefficients. This process



**Figure 5:** (a) Embedded image encoder and (b) embedded image decoder.

is illustrated by the block diagram shown in Fig. 5b. Since the arithmetic coding model is backward adaptive, we need not transmit it as side information. Furthermore, because the decoder's knowledge exactly mirrors that of the encoder at any given point in the processing of the bit stream, there is no need to transmit pass delimiters or synchronization signals (although these might be useful for resolution-scaleable compression). Also, the resolution enhancement bits transmitted during the subordinate pass do not need any location specifiers-- the decoder knows the exact transmission order of these bits because it has reconstructed the same subordinate list as the encoder had at that point in the processing.

## 4. Feature-Based Compression


### 4.1 Feature-selective resolution control

The conventional EZW algorithm allocates resolution uniformly across the image. To achieve this spatial uniformity, it actually distributes resolution to the wavelet coefficients non-uniformly across the wavelet scales (i.e., frequency subbands). Specifically, a

coarser scale is allocated twice as much resolution as the next finer scale; this allocation is implicitly controlled by the use of a unitary or unitary-like scaling in the wavelet decomposition. Such coefficient scaling has the effect of increasing the gain in each successive level of the 2D wavelet decomposition by a factor of 2. To understand how a multiplicative factor implicitly control resolution, consider the following example: assume that the true value of a wavelet coefficient is 85 and that the final dominant pass through the coefficients ends with threshold  $T = 64$ . Without rescaling, the final uncertainty interval for this coefficient in the decoder will be  $[64, 128)$ , resulting in a reconstructed coefficient value of  $96 \pm 32$ . Now, assume instead that this coefficient is multiplied by 2 prior to coding (i.e., we code the value 170). During the pass when  $T=128$ , the coefficient will be declared significant and approximated in the decoder by  $192 \pm 64$ . After the refinement pass, however, the new approximation will be  $160 \pm 32$ . Since the encoder stops after the dominant pass for  $T=64$ , the coefficient will receive no further refinement bits. Dividing the coefficient approximation by 2 restores its original scaling and results in the final

estimate of  $80 \pm 16$ . Thus, the uncertainty region of the new estimate, [64, 96], is 1/2 that of the original!

While scaling is used in the classical EZW algorithm to implicitly control the bit allocations to coefficients in the different wavelet scales, it can also be used explicitly to weight 'features' in the imagery. What is a feature? By our definition, it is anything in the imagery that can be localized in space and/or frequency. For example, a tank in a reconnaissance photo might be a spatial feature to which we want to allocate additional resolution. On the other hand, an orchard in the same photo could represent a space-frequency feature (i.e., a feature that is defined by a frequency spectrum within a spatial region) whose resolution allocation should be reduced in order to increase the resolution of other, more interesting portions of the image. Either way, the allocation of resolution (and thus bits) is most easily controlled by scaling coefficients up or down by the appropriate powers of 2 prior to encoding. It is important to note, however, that rescaling only adjusts the resolution of coefficients relative to other coefficients--this process is a zero sum game! Thus, scaling all of the wavelet coefficients up by a factor of 2 will have no effect on the resolution of the reconstructed image.




**Figure 6:** All coefficients shaded in gray are rescaled together.

## 4.2 Coding Scheme

By using coefficient rescaling, we create a single embedded bit stream in which different wavelet coefficients are represented with varying precision. If all of the coefficients corresponding to a single zerotree are coded at a specified precision, then the corresponding region of the reconstructed image will be reconstructed at the

same precision (assuming that orthonormal wavelets are used). The x's in Fig. 6 correspond to the coefficients which form a complete zerotree, and it is these coefficients which must be multiplied by a power of 2 scaling factor to increase the resolution of the corresponding 16x16 image region. The size of the region in the image which corresponds to 1 zerotree (the minimally controllable region) for a depth L wavelet decomposition is  $2^L \times 2^L$ . Thus, to *spatially* vary the resolution across the image, one need only varying the coefficient scaling on a zerotree by zerotree basis. If, on the other hand, one wishes to increase or decrease the resolution of specific frequency bands within a given region, one must vary the scaling factor between wavelet scales corresponding to the same zerotree. Note that by using a fixed wavelet decomposition, we have limited the amount of frequency control which can be achieved-- i.e., the bandwidth of each image subband increases by a factor of 2 in each dimension as the wavelet scale decreases. If one instead used an adaptive wavelet packet decomposition [12], [13], one could also optimize the bandwidth of the subbands for the space-frequency features of interest. Because such wavelet packets greatly increase the complexity of the encoder, however, we have chosen not to use them in our system.



**Figure 7:** Organization of compressed bit stream.

For the decoder to correctly reconstruct the wavelet coefficients, it must know which areas have been rescaled and by what scaling factor. The only exception to this general rule is when a coefficient has been scaled down so much that it will be reconstructed as 0 by the decoder. In the case where small areas of enhanced resolution are desired, the side information describing the rescaling to the decoder is very compact and does not require lossless compression. Figure 7 shows the organization of the compressed bit stream. First, the image mean (which was subtracted from the image prior to the wavelet decomposition) and starting threshold are sent; this is always done for any embedded compression algorithm. Next, a set of quality factors are transmitted, one for each zerotree whose resolution has been in-

increased. In the current instantiation, each quality factor is a 2 bit quantity and allows for three levels of resolution increase (x2, x4, and x8). Since the number of zerotrees with resolution increases is not predetermined, a fourth parsing symbol is also allowed which terminates the 'quality factors' section and tells the decoder how many 'ztr locations' to expect. Each 'ztr location' is a numeric value which uniquely indexes 1 zerotree; if the image is of size 512x512 and a depth 5 decomposition is used, then 1 byte is sufficient to uniquely identify a zerotree. One zerotree location is transmitted for each 'quality factor', and the ordering is such that each zerotree location index corresponds exactly to a previously transmitted quality factor. In addition, no parsing symbol is needed because the decoder already knows how many zerotree location indices to expect. Finally, the embedded bit stream containing the compressed representation of the actual image is transmitted to the decoder.

One might note that the proposed method of transmitting the side information will not be particularly efficient if a large number of zerotrees are rescaled. In such a situation, it might be more efficient to simply send 1 bit for each zerotree indicating whether or not its scaling has been altered along with an ordered list containing the magnitudes of the rescaling. In the case where there are 256 total zerotrees, transmitting the ztr location specifiers in this way would require exactly 32 bytes. Thus, the tradeoff is clear: if fewer than 32 zerotrees are rescaled, an explicit index to each should be transmitted; otherwise, the 1 bit/zerotree rescaling map should be sent.

## 5. Results

For the results presented in this section, we use a 5-level decomposition based on a 5/3 biorthogonal wavelet transform (called (2,2) in [14]). To increase its speed, we use lifting to implement this wavelet which allows the high and lowpass filters to share computations [15]. In addition, lifted transforms are in-place and thus do not require large amounts of scratch memory during computation. Since a 5-level decomposition has been selected, the minimally controllable spatial region is 32x32 (not taking into account the overlapping basis functions of the transform). From a sequence of 800 frames, a training set of ten images containing the group of four buildings indicated by the white arrow in Fig. 8a. was selected.

From this training set, a MACH filter was constructed, to recognize the buildings as the feature of interest. The size of the filter in this case was 64x64 which is larger than the minimally controllable region

as determined by the depth of the wavelet transform. The regions with the top 3 correlation peaks (i.e., the best 3 matches) are allocated resolution according to the following formula: #1 always receives 8 times the resolution of the background; #2 and #3 receive the same if their correlation peaks are within 5% of #1, but they receive x4 resolution if the peaks are within 20%, and x2 resolution if less than 20%. The reasoning behind such an allocation scheme is that while the MACH filter ideally will always return the highest correlation peaks for the true target, this is not always the case. If a false alarm has the highest peak value, presumably the true target has a similar value and hence correlation peaks within 5% of the top peak receive the same resolution. If the correlation scores are significantly different, the probability that the lower scores are true targets decreases and the region receives a lower bit allocation. Figures 8 and 9 illustrate the results for compression at ratios of 80:1 and 160:1, respectively. By looking at the error residuals (Figures 8b and 9b), one can see the differences in the allocation of resolution to the different regions. In both examples, the area containing the group of four buildings has the highest resolution while the other two highlighted regions have lower resolution (but still higher than the background). As mentioned previously, the choice of three high quality regions is somewhat arbitrary but it does depend some on the design of the MACH filter. A tradeoff between the amount of compression and the probability of detection and false alarms needs to be considered when actually implementing this algorithm in practice.

## 6. Discussion and Conclusions

In the previous sections we have described a synthesis of two disparate areas of image processing: compression and ATR/ATC. The regional compression algorithm expands the effective bandwidth available by maximizing the information content of the transmitted imagery, using information from the ATR algorithm. One obvious area where this technique may prove valuable is in the UAV/UCAV arena. For example, with the use of the ATR/ATC driven compression, analysis by a human operator is simplified, rapidly identifying potential areas of interest. A second area where this work could be extended to is image database management. By utilizing regional compression driven by an ATR/ATC algorithm it is possible to achieve useful compression ratios (>80:1) while preserving image quality in areas that may be of interest. The potential savings in communications bandwidth could enable more vehicles to operate simul-


aneously than would be possible with standard compression techniques in use currently.

## References

- [1] D. Carlson, "Optimal tradeoff composite correlation filters," Ph.D. thesis Carnegie Mellon Univ., Oct. 1996.
- [2] A. Mahalanobis *et al*, "Unconstrained correlation filters," *Applied Optics*, vol. 33, pp. 3751-3759, 1994.
- [3] A. Mahalanobis, B.V.K. Vijaya Kumar, and S.R.F. Sims, "Distance classifier correlation filters for multiclass target recognition," *Applied Optics*, vol. 35, pp. 3127-3133, 1996.
- [4] B.V.K. Vijaya Kumar, D. Carlson, and A. Mahalanobis, "Optimal tradeoff synthetic discriminant function (OTSDF) filters for arbitrary devices," *Optics Letters*, vol. 19, pp. 1556-1558, 1994.
- [5] Ph. Refregier, "Optimal tradeoff filters for noise robustness, sharpness of the correlation peak and Horner efficiency," *Optics Letters*, vol. 16, pp. 829-831, 1991.
- [6] Z. Bahri and B.V.K. Vijaya Kumar, "Generalized synthetic discriminant functions," *J. Opt. Soc. Am. A*, vol. 5, pp. 562-571, 1988.
- [7] B.V.K. Vijaya Kumar, "Minimum variance synthetic discriminant functions," *J. Opt. Soc. Am. A*, vol. 3, pp. 1574-1584, 1986.
- [8] A. Mahalanobis, B.V.K. Vijaya Kumar, and D. Casasent, "Minimum average correlation energy filters," *Applied Optics*, vol. 26, pp. 3633-40, 1987.
- [9] J.M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. on Signal Processing*, Vol. 41, No. 12, pp. 3445-3462, December 1993.
- [10] J.M. Shapiro, "A fast technique for identifying zeros in the EZW algorithm," *Proc. IEEE Int. Conf. on Acoust., Speech, and Signal Proc.*, Atlanta, GA, May 1996.
- [11] I.H. Witten, R.M. Neal, and J.G. Cleary, "Arithmetic coding for data compression," *Communications of the ACM*, Vol. 30, No. 6, June 1987, pp. 520-540.
- [12] R.R. Coifman and M.V. Wickerhauser, "Entropy-based algorithm for best basis selection," *IEEE Trans. on Information Theory*, Vol. 38, No. 2, pp. 713-18, March 1992.
- [13] K. Ramachandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Trans. on Image Processing*, Vol. 2, No. 2, April 1993, pp. 160-175.
- [14] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, PA, 1992.
- [15] W. Sweldens, "The lifting scheme: a new philosophy in biorthogonal wavelet constructions," *Proceedings of the SPIE (Wavelet Applications in Signal Proc. III)*, Vol. 2569, 1995, pp. 68-79.




(a)




(b)

**Figure 8:** (a) Reconstruction of image compressed by 80:1 ratio. Squares have been added to highlight enhanced regions and arrow indicates the particular region of interest for which the system was trained. (b) Error residual between reconstructed and original image where white areas in residual denote large errors.





(a)



(b)

**Figure 9:** (a) Reconstruction of image compressed by 160:1 ratio. Again, squares have been added to highlight enhanced regions. (b) Error residual between reconstructed and original image where white areas in residual denote large errors.